Storage of quantum dynamics in quantum states:
a quasi-perfect programmable quantum gate

G. Vidal and J. I. Cirac

Institut für Theoretische Physik, Universität Innsbruck,A-6020 Innsbruck, Austria

(February 1, 2008)

We show how quantum dynamics (a unitary transformation) can be captured in the state of a quantum system, in such a way that the system can be used to perform, at a later time, the stored transformation almost perfectly on some other quantum system. Thus programmable quantum gates for quantum information processing are feasible if some small degree of imperfection is allowed. We discuss the possibility of using this fact for securely computing a secret function on a public quantum computer. Finally, our scheme for storage of operations also allows for a new form of quantum remote control.

PACS Nos. 03.67.-a, 03.65.Bz

I. INTRODUCTION

Quantum information theory explores the potential of quantum mechanics in order to process and transmit information. A two-level system, a qubit, constitutes the unit resource for storing information. Similarly, a unitary operation on one qubit can be regarded as a basic unit of information processing [1]. In this paper we explore the possibility of storing quantum dynamics, in particular unitary transformations, in the state of a quantum system, in a manner that the transformation can be performed at a later time and on another system almost perfectly.

A. Quantum programmable gates.

The problem we address can be well-posed in the context of quantum circuitry. We will say that the *program state* $\mathcal{P}_U$ of some *program register* stores the one-qubit transformation U, if some "fixed" protocol employing the state $\mathcal{P}_U$ is able to perform U on an arbitrary *data* state ρ of a single qubit *data register*. Here, a "fixed" protocol means that the manipulation of the joint state

$$\rho \otimes \mathcal{P}_U \qquad (1)$$

does not require knowing the operation U nor the state ρ. A device able to transform state (1) into

$$U\rho U^{\dagger} \otimes \mathcal{R}_{\rho,U}, \qquad (2)$$

where $\mathcal{R}_{\rho,U}$ is just some residual state, is known as a *programmable* quantum gate [2]. In a similar fashion as modern (classical) computers take both the program to be executed and the data to be processed as input bit strings, a programmable or universal quantum gate is a device whose action U on an arbitrary data state ρ is determined by the program state $\mathcal{P}_U$.

Nielsen and Chuang [2] analyzed the possibility of constructing such a programmable quantum gate. Its total dynamics was described by means of a fixed unitary operator G according to

$$G[|d\rangle \otimes |\mathcal{P}_U\rangle] = (U|d\rangle) \otimes |\mathcal{R}_U\rangle, \qquad (3)$$

where only pure data states $|d\rangle$ were considered because this already warranties the mixed state case. Notice that the program state $|\mathcal{P}_U\rangle$ and the residual state $|\mathcal{R}_U\rangle$ — which was showed to be independent of $|d\rangle$— can always be taken to be pure, by extending the program register with an ancillary system if needed. Nielsen and Chuang proved that any two inequivalent operations U and V require orthogonal program states, that is $\langle \mathcal{P}_U|\mathcal{P}_V\rangle = 0$. Thus, in order to perfectly store a given operation U_i from some set $\{U_i\}_{i\in I}$, a vector state $|\mathcal{P}_{U_i}\rangle$ from an orthonormal basis $\{|\mathcal{P}_{U_i}\rangle\}_{i\in I}$ has to be used. The operation U_i can then be implemented by, say, measuring the program register to obtain the value i, and gauging correspondingly some convenient experimental device. Since the set of unitary operations is infinite, their result implied that no universal gate can be constructed using finite resources, that is, with a finite dimensional program register.

B. Main results.

The aim of this work is to present programmable quantum gates with a finite program register, and thus physically feasible. A finite register turns out to be sufficient if a degree of imperfection, no matter how small, is allowed in performing the unkwon operation U. We will construct a family of *probabilistic* programmable quantum gates, that is programmable quantum gates which work with a given prior probability $p \geq 1 - \epsilon$ of a successful implementation of U. Such a one-qubit gate with $\epsilon = 3/4$ was already described in [2]. Here we will achieve any arbitrarily small $\epsilon > 0$. We will also consider *approximate* programmable quantum gates, which perform

an operation $\mathcal{E}_U$ very similar to the desired U, that is $F(\mathcal{E}_U, U) \geq 1 - \epsilon$ for some transformation fidelity F.

The second main result is a lower bound on the dimension of the program register of the programmable gate in terms of its degree of imperfection ϵ. It implies that the orthogonality result of [2] is robust. We will discuss its implications in the context of secure secret computation.

Finally, operations stored in a quantum state can be teleported. This leads to a new scheme for quantum remote control [5] that only requires unidirectional communication.

II. QUASI-PERFECT PROGRAMMABLE QUANTUM GATES

We start by showing how to store and reimplement, in an imperfect but feasible fashion, an arbitrary one-qubit unitary operation of the form

$$U_\alpha \equiv \exp(i\alpha\sigma_z), \quad (4)$$

where $\alpha \in [0, \pi)$. Notice that a general one-qubit operation $U \in SU(2)$ can be obtained by composing three operations of the form of eq. (4) with some fixed unitary operations, for instance as $U = U_{\alpha_3}\exp(-i\pi\sigma_x/2)U_{\alpha_2}\exp(i\pi\sigma_x/2)U_{\alpha_1}$.

A. Single-qubit program state.

Let us consider the state

$$|\alpha\rangle \equiv \frac{1}{\sqrt{2}}(e^{i\alpha}|0\rangle + e^{-i\alpha}|1\rangle), \quad (5)$$

which someone, say Alice, can prepare by applying U_α on a qubit in the standard state $(|0\rangle + |1\rangle)\sqrt{2}$. Suppose she also prepares, along with $|\alpha\rangle$, another qubit in some arbitrary state $|d\rangle = a|0\rangle + b|1\rangle$ and provides Bob, who doesn't know α nor the complex coefficients a and b, with the two qubits in state $|d\rangle \otimes |\alpha\rangle$. Alice challenges now Bob to obtain the state $U_\alpha|d\rangle$.

What Bob can do in order to implement the unknown U with some probability of success is to perform a C-NOT operation taking the data qubit in state $|d\rangle$ as the control and the program qubit in state $|\alpha\rangle$ as the target. This will constitute the basic part of our simplest programmable quantum gate. Recalling that the C-NOT gate,

$$|0\rangle\langle0| \otimes I + |1\rangle\langle1| \otimes \sigma_x, \quad (6)$$

permutes the $|0\rangle$ and $|1\rangle$ states of the target (second qubit) only if the control (first qubit) is in state $|1\rangle$, it is easy to check that the two-qubit state is transformed according to

$$|d\rangle \otimes |\alpha\rangle \xrightarrow{\text{C-NOT}} \frac{1}{\sqrt{2}}(U_\alpha|d\rangle \otimes |0\rangle + U_\alpha^\dagger|d\rangle \otimes |1\rangle). \quad (7)$$

Therefore, a projective measurement in the $\{|0\rangle, |1\rangle\}$ basis of the program register will make the data qubit collapse either into the desired state $U_\alpha|d\rangle$ or into the wrong state $U_\alpha^\dagger|d\rangle$, with each outcome having prior probability $1/2$. That is, we have already constructed a probabilistic programmable quantum gate with error rate $\epsilon = 1/2$ (see figure (1)). Notice that a single qubit has been sufficient for Alice to store an arbitrary unitary U_α, i.e., one from an infinite set, although its recovery only succeeds with probability $1/2$. If Bob obtains $U_\alpha^\dagger|d\rangle$ instead of $U_\alpha|d\rangle$, then not only he fails at performing the wished operation, but in addition he does no longer have the initial data state $|d\rangle$.

B. Multi-qubit programs.

How can we construct a more efficient programmable gate? Notice that in case of failure, a second go of the previous gate can correct $U_\alpha^\dagger|d\rangle$ into $U_\alpha|d\rangle$. Indeed, Bob needs only apply the gate of fig. (1) to $U_\alpha^\dagger|d\rangle$, inserting a new program state, namely $|2\alpha\rangle$, which Alice can prepare by performing twice the operation U_α on $(|0\rangle + |1\rangle)/\sqrt{2}$. Therefore, if Alice supplies the state $|\alpha\rangle \otimes |2\alpha\rangle$ to Bob, he can perform the operation U_α with probability $3/4$. Figure (2) displays a more compact version of this second probabilistic programmable gate, which requires a two-qubit program register and has a probability of failure $\epsilon = 1/4$.

In case of a new failure, the state of the system becomes $U_\alpha^{\dagger 3}|d\rangle$. Bob can insert again this state, together with state $|4\alpha\rangle$, into the elementary gate. If Bob has no luck and keeps on obtaining failures, he can try to correct the state as many times as he wishes, provided that the state $|2^l\alpha\rangle$ is available for the lth attempt. Therefore, for any N, the N-qubit state $\otimes_{l=1}^N |2^l\alpha\rangle$ can be used to implement the transformation U_α with probability $1 - (1/2)^N$.[1] The corresponding probabilistic programmable gate (see figure (3)), consists of the unitary transformation of $|d\rangle \otimes (\otimes_{l=1}^N |2^l\alpha\rangle)$ into

$$\frac{1}{2^{N/2}}(\sqrt{2^N - 1}\, U_\alpha|d\rangle \otimes |\text{r}\rangle + U_\alpha^{(2^N-1)\dagger}|d\rangle \otimes |\text{w}\rangle) \quad (8)$$

---

[1] Note that our several-step correcting scheme for implementing U_α resembles that used in [7] to implement a nonlocal unitary operation. In the present context all intermediate measurements and conditional actions can be substituted by a single unitary operation, as described in Figures 2 and 3. In this section we have first presented the several-measurement version for pedagogical reasons.

and of a posterior measurement of the program register (either in state $|r\rangle$ or $|w\rangle$, $\langle r|w\rangle = 0$). Its failure probability, $\epsilon = (1/2)^N$, decreases exponentially with the size N of the program register.

It is interesting to look at how long the program needs to be, on average, until Bob succeeds to perform U_α with certainty. With probability $p_1 = 1/2$ he succeeds after using a single-qubit program; with probability $p_2 = 1/4$ a two-qubit program is sufficient; etc. The average length $\bar{N}$ of the required program is thus

$$\bar{N} = \sum_{N=1}^{\infty} p_N N = \sum_{N=1}^{\infty} \frac{N}{2^N} = 2. \qquad (9)$$

That is, a two qubit register is sufficient, on average, to store an arbitrary U_α so that it can be performed with certainty.

C. Probabilistic versus approximate programmable gates.

A probabilistic programmable gate may either succeed or fail, depending on the result of the final measurement on the program register. An approximate gate, instead, performs a transformation only similar to the desired one, but it is always successful. Suppose we want to apply the unitary transformation U on $|\psi\rangle$ but instead another (general) transformation $\mathcal{E}$ is actually performed. A possible way of quantifying how similar these two operations are is by applying both operations to the same state $|\psi\rangle$, and then computing the fidelity between the two transformed states, $\mathcal{E}(|\psi\rangle\langle\psi|)$ and $U|\psi\rangle$. When averaged over all possible $|\psi\rangle$ this reads

$$F(\mathcal{E}, U) \equiv \int d\psi \langle\psi|U^\dagger \mathcal{E}(|\psi\rangle\langle\psi|)U|\psi\rangle. \qquad (10)$$

Suppose now that after the transformation (8) of the previous probabilistic gate we decide to ignore the state of the N-qubit program register. Then the programmable gate works approximately, implementing an operation $\mathcal{E}_U(\rho) \equiv (1-\epsilon)U_\alpha \rho U_\alpha^\dagger + \epsilon \tilde{U}\rho\tilde{U}^\dagger$, where $\tilde{U} \equiv U_{(1-2^N)\alpha}$. The average fidelity of performance (10) satisfies $F(\mathcal{E}_U, U) \geq 1 - \epsilon = 1 - (1/2)^N$.

III. PROGRAMABLE GATES AND SECRET COMPUTATION ON A PUBLIC COMPUTER

So far we have explicitly constructed programmable quantum gates that perform, either probabilistically or approximately, some class of one-qubit unitary operations U_α. But the previous protocols also allow Alice to codify with finite resources any unitary operation $V^{(l)}$ acting on an arbitrary number l of qubits. Indeed, as already mentioned, Alice can codify an arbitrary one-qubit

unitary operation using only 3 U_α's, and then also combine several of those with C-NOT gates to obtain $V^{(l)}$.

A. Secret computation on a public computer?

In view of these results, one may wonder whether quasi-perfect programmable gates can be applied, in the context of quantum cryptography and computation, to secretly compute some unitary operation $V^{(l)}$, for instance a precious algorithm, on some initial l-qubit state $|d^{(l)}\rangle$. The idea is that Alice gives a program state $|\mathcal{P}_{V^{(l)}}\rangle$ and the data state $|d^{(l)}\rangle$ to Bob, who operates a programmable quantum gate array but ignores $V^{(l)}$. Bob is required to compute $V^{(l)}|d^{(l)}\rangle$, but Alice does not want Bob to know what program he is running on his quantum computer. If the gate is perfect as in (3), then Bob can in principle distinguish $|\mathcal{P}_{V^{(l)}}\rangle$ from any other program state, since they are orthogonal. Therefore he can, imperceptibly to Alice, make an illegal copy of the program, perform the required transformation using the original program state, and give the computed state to Alice.

However, when the gate is slightly imperfect, different programs need no longer be orthogonal. Now Bob can not determine perfectly well which program he is to run in his computer. If he tries to estimate $|\mathcal{P}_{V^{(l)}}\rangle$, then in addition he will necessarily modify the program state, which will result in an improper performance of the gate and then Alice—who may have simply been testing Bob's integrity—can, in principle, detect it. That is, it is not possible for Bob to copy, even in an approximate form, the program state and at the same time perform the operation Alice has commended him with, without this being detectable.

We next derive lower bounds on the size N of the program register of any quasi-perfect (i.e. with $\epsilon \to 0$) programmable gate, and on the degree of orthogonality between its program states $|\mathcal{P}_U\rangle$ corresponding to similar operations, in terms of its failure parameter ϵ. These bounds represent a severe limitation on the degree of reliability that a security scheme based on the above ideas can offer. They indicate that the program vectors $|\mathcal{P}_U\rangle$ and $|\mathcal{P}_V\rangle$ are significantly non-orthogonal (that is, non-distinguishable) only when the imperfection parameter ϵ makes them effectively equivalent.

B. Upper bound to the indistinguishability of states for different programs.

Let us consider a generic imperfect programmable gate acting on a $\mathcal{C}^n$ system, so that it can be programmed to perform some or all $U \in SU(n)$. It can be described by a unitary operator G_ϵ according to

$$G_\epsilon[|d\rangle \otimes |\mathcal{P}_U\rangle] = \sqrt{1-\epsilon}(U|d\rangle) \otimes |\mathcal{R}_U^d\rangle + \sqrt{\epsilon}|w_U^d\rangle, \quad (11)$$

where the wrong state $|w_U^d\rangle$ is not required to fulfill any requirement for an *approximate* gate, whereas it must satisfy $\langle\mathcal{R}_V^{d_1}|w_U^{d_2}\rangle = 0$ (here 0 is the null vector of the data register) for any two inequivalent operations V and U and any two data states $|d_1\rangle$ and $|d_2\rangle$ for a *probabilistic* gate. This last condition is necessary for Bob, who ignores both U and the data, to be able to know whether the transformation U has been successfully performed by measuring the program register.

We first notice that the state $|\mathcal{R}_U^d\rangle$ only depends on d through a contribution of order ϵ^τ, where from now on $\tau = 1/2$ for approximate gates and $\tau = 1$ for probabilistic ones. Indeed, for any program state $|\mathcal{P}_U\rangle$, the scalar product of (11) corresponding to any two data states $|d_1\rangle$ and $|d_2\rangle$ reads

$$\langle d_1|d_2\rangle = \langle d_1|d_2\rangle\langle\mathcal{R}_U^{d_1}|\mathcal{R}_U^{d_2}\rangle + O(\epsilon^p), \qquad (12)$$

$[O(\epsilon^\tau)$ is a term linear in $\epsilon^\tau]$ from which, by fixing $|d_1\rangle$ and considering any $|d_2\rangle$, we find that $\langle\mathcal{R}_U^{d_1}|\mathcal{R}_U^{d_2}\rangle = 1 + O(\epsilon^\tau)$. That is

$$|\mathcal{R}_U^d\rangle = |\mathcal{R}_U\rangle + O(\epsilon^\tau). \qquad (13)$$

Keeping this in mind, we now consider, for any given $|d\rangle$, the scalar product of (11) corresponding to two unitary operations U and V, which turns out to read

$$\langle\mathcal{P}_U|\mathcal{P}_V\rangle = \langle d|U^\dagger V|d\rangle\langle R_U|R_V\rangle + O(\epsilon^\tau). \qquad (14)$$

The scalar product $\langle\mathcal{P}_U|\mathcal{P}_V\rangle$ does not depend on $|d\rangle$. Therefore the dependence of $\langle d|U^\dagger V|d\rangle\langle R_U|R_V\rangle$ on $|d\rangle$ has to be of order ϵ^τ, at most. Suppose U and V are very close. That is,

$$U^\dagger V = e^{iL} = I + iL - \frac{1}{2}L^2 + O(L^3), \qquad (15)$$

where all the eigenvalues $\alpha_1 \geq \alpha_2 \geq ... \geq \alpha_n$ of the traceless $(\sum_i \alpha_i = 0)$ hermitian operator $L = \sum_i^n \alpha_i|i\rangle\langle i|$, are very small. The largest variation of $\langle d|U^\dagger V|d\rangle$ in (14) for two different vectors $|d\rangle$ is $\langle 1|U^\dagger V|1\rangle - \langle n|U^\dagger V|n\rangle = \alpha_1 + |\alpha_n|$. We introduce a distance on the set of operators on C^n,

$$D_{U,V} \equiv (\text{Tr}[(U-V)^\dagger(U-V)])^{\frac{1}{2}}. \qquad (16)$$

Then $\alpha_1 + |\alpha_n| \geq (\sum_i \alpha_i^2)^{1/2}/n = D_{U,V}/\sqrt{2}n$. Subtracting (14) for $|d\rangle = |n\rangle$ from itself for $|d\rangle = |1\rangle$ we conclude that $|\langle\mathcal{R}_U|\mathcal{R}_V\rangle| \leq O(\epsilon^\tau)n/D_{U,V}$, which finally implies

$$|\langle\mathcal{P}_U|\mathcal{P}_V\rangle| \leq \frac{O(\epsilon^p)n}{D_{U,V}}. \qquad (17)$$

This bound says that in a programmable quantum gate with a small error rate $\epsilon \ll 1$, two transformations U and $V \in SU(n)$ will have program states with significant overlap $\langle\mathcal{P}_U|\mathcal{P}_V\rangle$ (states $|\mathcal{P}_U\rangle$ and $|\mathcal{P}_V\rangle$ are indistinguishable) only if U and V are also very close to each other, $D_{U,V} \ll 1$. That is, only if U and V process the data very similarly, then a dishonest Bob is unable to distinguish between the corresponding programs.

C. Lower bound to the dimension of the program register.

The previous result can also be used to derive a lower bound on the dimension of the program register of an imperfect programmable gate with error ϵ. For simplicity we will assume that the gate can only be programmed to perform the one-qubit transformations U_α from eq. (4). Consider a discrete subset of such transformations, namely those with $\alpha_s \equiv \pi s/M$, $s = 0, ..., M-1$, and apply the previous bound to U_{α_s} and $U_{\alpha_{s+1}}$. We obtain

$$|\langle\mathcal{P}_{\alpha_s}|\mathcal{P}_{\alpha_{s+t}}\rangle| \leq K\epsilon^p M, \quad \forall t \neq 0, \qquad (18)$$

where K is some unimportant constant. We need the following lemma.

Lemma: Let $\{|\psi_i\rangle\}_{i=1}^q$ be a set of q (normalized) vectors such that their scalar products $\nu_{ij} \equiv \langle\psi_i|\psi_j\rangle$ satisfy $|\nu_{ij}| < q^{-1}$ for $i \neq j$. Then the q vectors are linearly independent.

Proof: The rank of the set $\{|\psi_i\rangle\}_{i=1}^q$ is equal to the rank of the matrix N, $N_{ij} \equiv \langle\psi_i|\psi_j\rangle$, which has ones in all diagonal entries. The modulus of any entry of the matrix $N-I$ is smaller than q^{-1}. Let $|\varphi\rangle$ be a normalized eigenvector of $N-I$, with eigenvalue λ. Then $\lambda[|\varphi\rangle]_i = [(N-I)|\varphi\rangle]_i = \sum_{j=1}^q \nu_{ij}[|\varphi\rangle]_j$, where $[\ ]_j$ denotes the jth vector component. Let i be such that $|[|\varphi\rangle]_i| \geq |[|\varphi\rangle]_j|$ $\forall j$. Then $|\sum_{j=1}^q \nu_{ij}[|\varphi\rangle]_j| \leq \sum_{j=1}^q |\nu_{ij}||[|\varphi\rangle]_i| < |[|\varphi\rangle]_i|$, that is, $|\lambda| < 1$, and since this holds for all the eigenvalues of $N-I$, N has q positive eigenvalues or, equivalently, rank q.

Remark: For q sufficiently large, if all ν_{ij} are of order ν and the components of the eigenvector $|\varphi\rangle$ are relatively equally weighted, then it is plausible that the eigenvalue $\lambda = \sum_{j=1}^q \nu_{ij}[|\varphi\rangle]_j/[|\varphi\rangle]_i$ is of the order $\nu q^{1/2}$ (random walk). This suggests that in order for the set $\{|\psi_i\rangle\}_{i=1}^q$ to have rank close to q, it is sufficient that $|\nu_{ij}| < q^{-1/2}$, instead of $|\nu_{ij}| < q^{-1}$ as required in the lemma.

Let us set $M \equiv (\epsilon^\tau K)^{-1/2}$. Then (18) becomes $|\langle\mathcal{P}_{\alpha_s}|\mathcal{P}_{\alpha_{s+t}}\rangle| \leq 1/M$, and this means, because of the lemma, that at least an M-dimensional Hilbert is required to contain $\{|\mathcal{P}_{\alpha_s}\rangle\}_{s=0}^{M-1}$. That is, the program register must consists of at least $(\tau/2)\log(1/\epsilon)$ qubits. Notice that the previous remark suggests that this bound may be reduced to $\tau\log(1/\epsilon)$ qubits, in which case the probabilistic programmable gate of figure (3) would require, asymptotically, the smallest possible program register. For a general programmable gate implementing some or all transformations $U \in SU(n)$ it is straightforward to obtain a similar lower bound on the dimensions of the program register, which also says that its number N of qubits grows proportionally to the logarithm of the inverse of the rate error, $N = k_n\log(1/\epsilon)$, for some positive constant k_n.

IV. APPLICATIONS: MANIPULATION OF UNKNOWN QUANTUM DYNAMICS

We have shown how to store an arbitrary unitary transformation in the pure state of a finite quantum register, in such a way that it can be performed quasi-perfectly at a later time. Once the unknown operation has been encoded in a quantum state, it can of course be processed using *any* known state manipulation technique.

A. Unidirectional quantum remote control.

An interesting application of our results is in the context of quantum remote control. As introduced by Huelga *et al.* in [5], let us suppose Bob wishes to manipulate some data state according to an unknown operation Alice, a distant party, can implement by using some device. If the state of Alice's device cannot be teleported, then the optimal protocol [5] is to use standard teleportation [4] to send the data from Bob to Alice, who will use the device to process it and will teleport it back to Bob.

But we now know how to efficiently store operations in quantum states, which can then be teleported. This leads to a new scheme for quantum remote control: Alice stores the operation in a quantum state and applies standard teleportation to send it to Bob.

Remarkably enough, in this protocol only one-way communication is required —in addition to entanglement —in order for Alice to remotely manipulate Bob's data, as opposed to the two-way classical communication of the scheme presented in [5]. This implies that the operation can be teleported independently of whether Bob's data state is already available. More specifically, we find that a N-qubit program $|\alpha\rangle \otimes ... \otimes |2^N \alpha\rangle$ can be teleported from Alice to Bob by using up N ebits of entanglement and by sending N classical bits from Alice to Bob (recall that the classical communication cost of quantum teleportation of equatorial states, as $|\alpha\rangle$, require only one bit per state [6]). Eq. (9) implies that, on average, 2 ebits of entanglement between Alice and Bob, and 2 classical bits from Alice to Bob are sufficient for Alice to teleport an arbitrary U_α to Bob, so that he, ignoring U_α, can perform it with certainty.

B. Estimation of quantum dynamics and storage of non-local transformations.

The storage of quantum transformations turns out to be useful in several other contexts. If state estimation techniques are applied to the system that stores an unknown operation, then we obtain the scheme for estimation of quantum dynamics recently exploited by Acín *et al* [3].

Cirac *et al* [7] have recently explored the possibilities of encoding operations in quantum states in the context of non-local transformations of a composite system. In particular, they have shown how to implement non-local unitary transformations using less than one ebit of entanglement. In an extension of their work, Dür *et al* [8] have considered alternative schemes for storing and manipulating quantum transformations.

V. CONCLUSIONS

We have presented a scheme for storing unitary operations in the quantum state of a finite dimensional program register. The operations can be implemented at a later time with some associated error ϵ, which decreases exponentially with the number of qubits of the program register. We have presented both probabilistic and approximate programmable quantum gates, and have discussed the possibility of using them to make a secrete computation on a public quantum computer. Finally, a unidirectional scheme for remote manipulation of quantum states has also been put forward.

ACKNOWLEDGMENTS

[1] Indeed, an arbitrary unitary operation $V^{(l)}$ on l qubits can decomposed as a fixed —that is, $V^{(l)}$-independent— sequence of ($V^{(l)}$-dependent) one-qubit unitary operations and a specific kind of two-qubit unitary operations, for instance C-NOT gates —since the sequence is fixed, $V^{(l)}$ is specified by the one-qubit transformations only; an arbitrary von Neumann measurement, given by orthogonal projectors $\{P_k\}$, can be decomposed into a general unitary operation V, with $V P_k V^\dagger \equiv Q_k$, a von Neumann measurement on a fixed set of projectors $\{Q_k\}$ and a second unitary operation V^{-1}. Finally, a general quantum operation on an arbitrary system can be achieved by appending an ancilla in some blank state to the system, performing a unitary transformation U followed by a von Neumann measurement, $\{P_k\}$, and disposing of part of the system and/or ancilla. Therefore any possible transformation of a

quantum system can be essentially specified by means of one-qubit unitary operations.

[2] M.A. Nielsen and I.L. Chuang, Phys. Rev. Lett. **79** 321 (1997).

[3] A. Acín, E. Jané and G. Vidal, "Optimal estimation of quantum dynamics", quant-ph/0012015.

[4] C. Bennett, G. Brassard, C. Crepeau, R. Josza, A. Peres and W.K. Wootters, Phys. Rev. Lett. 70, 1895 (1993).

[5] S.F. Huelga, J.A. Vaccaro, A. Chefles and M.B. Plenio, "Quantum remote control: Teleportation of unitary operations", quant-ph/0005061.

[6] See, for instance, H.-K. Lo, "Classical Communication Cost in Distributed Quantum Information Processing - A generalization of Quantum Communication Complexity", quant-ph/9912009.

[7] J. I. Cirac, W. Dür, B. Kraus, M. Lewenstein, "Entangling operations and their implementation using a small amount of entanglement", quant-ph/0007057.

[8] W. Dür and J. I. Cirac, in preparation.

FIG. 3. This probabilistic programmable quantum gate uses a N-qubit program register and succeeds with probability $p = 1 - (1/2)^N$, i.e. $\epsilon = (1/2)^N$. If no final measurement on the program register is made, or its result is ignored, then this circuit can be regarded as an approximate programmable quantum gate with performance fidelity $F \geq 1 - (1/2)^N$.
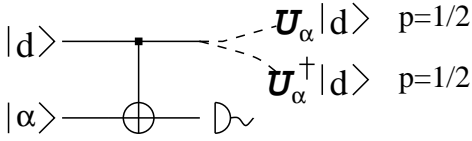

FIG. 1. This simple quantum circuit implements a probabilistic gate that takes unknown data and program states $|d\rangle$ and $|\alpha\rangle$ and produces, depending on the result of a measurement on the program register, either $U_\alpha|d\rangle$ or $U_\alpha^\dagger|d\rangle$, with equal prior probability $1/2$ (i.e. $\epsilon = 1/2$).
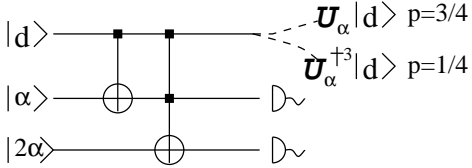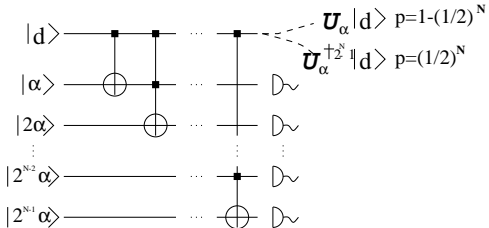

FIG. 2. The gate of figure (1) can be improved by making a conditional correction of the output after its C-NOT gate. This is achieved by means of a Toffoli gate, which acts as a C-NOT between the first and third line of the circuit only when the second line carries a $|1\rangle$, which corresponds to a failure in the circuit of figure (1). Bob can measure the second and third lines of the circuit. Only if he obtains 1 for both outcomes (which happens only one forth of the times, $\epsilon = 1/4$) is the transformation unsuccessful.